

AUTOCRATIC STRATEGIES FOR ITERATED GAMES WITH ARBITRARY ACTION SPACES

ALEX MCAVOY AND CHRISTOPH HAUERT

ABSTRACT. The recent discovery of zero-determinant strategies for the iterated Prisoner’s Dilemma sparked a surge of interest in the surprising fact that a player can exert unilateral control over iterated interactions. These remarkable strategies, however, are known to exist only in games in which players choose between two alternative actions such as “cooperate” and “defect.” Here we introduce a broader class of autocratic strategies by extending zero-determinant strategies to iterated games with more general action spaces. We use the continuous Donation Game as an example, which represents an instance of the Prisoner’s Dilemma that intuitively extends to a continuous range of cooperation levels. Surprisingly, despite the fact that the opponent has infinitely many donation levels from which to choose, a player can devise an autocratic strategy to enforce a linear relationship between his or her payoff and that of the opponent even when restricting his or her actions to merely two discrete levels of cooperation. In particular, a player can use such a strategy to extort an unfair share of the payoffs from the opponent. Therefore, although the action space of the continuous Donation Game dwarfs that of the classical Prisoner’s Dilemma, players can still devise relatively simple autocratic and, in particular, extortionate strategies.

1. INTRODUCTION

Game theory provides a powerful framework to study interactions between individuals (“players”). Among the most interesting types of interactions are social dilemmas, which result from conflicts of interest between individuals and groups (Dawes, 1980; Van Lange et al., 2013). Perhaps the most well-studied model of a social dilemma is the Prisoner’s Dilemma (Axelrod and Hamilton, 1981). A two-player game with actions, C (“cooperate”) and D (“defect”), and payoff matrix,

$$\begin{array}{cc} & \begin{matrix} C & D \end{matrix} \\ \begin{matrix} C \\ D \end{matrix} & \begin{pmatrix} R & S \\ T & P \end{pmatrix}, \end{array} \quad (1)$$

is said to be a Prisoner’s Dilemma if $T > R > P > S$ (Axelrod, 1984). In a Prisoner’s Dilemma, defection is the dominant action, yet the players can realize higher payoffs from mutual cooperation (R) than they can from mutual defection (P), resulting in a conflict of interest between the individual and the pair, which characterizes social dilemmas. Thus, in a one-shot game (i.e. a single encounter), two opponents have an incentive to defect against one another, but the outcome of mutual defection (the unique Nash equilibrium) is suboptimal for both players.

One proposed mechanism for the emergence of cooperation in games such as the Prisoner’s Dilemma is direct reciprocity (Trivers, 1971; Nowak, 2006b), which entails repeated encounters between players and allows for reciprocation of cooperative behaviors. In an iterated game, a player might forgo the temptation to defect in the present due to the threat of future retaliation—“the shadow of the future”—or the possibility of future rewards for cooperating (Axelrod, 1984; Delton et al., 2011), phenomena for which there is both theoretical and empirical support (Heide and Miner, 1992; Dal Bó, 2005). One example of a strategy for the iterated game is to copy the action of the opponent in the previous round (“tit-for-tat”) (Axelrod, 1984). Alternatively, a player might choose to retain his or her action from the previous round if and only if the most recent payoff was R or T (“win-stay, lose-shift”) (Nowak and Sigmund, 1993). These examples are among the simplest and most successful strategies for the iterated Prisoner’s Dilemma (Nowak, 2006a).

In a landmark paper, Press and Dyson (2012) deduce the existence of zero-determinant strategies, which allow a single player to exert much more control over this game than previously thought possible. Since their introduction, these strategies have been extended to cover multiplayer social dilemmas (Hilbe et al., 2014b; Pan et al., 2015) and temporally-discounted games (Hilbe et al., 2015a). Moreover, zero-determinant

strategies have been studied in the context of evolutionary game theory (Hilbe et al., 2013a; Adami and Hintze, 2013; Stewart and Plotkin, 2013; Szolnoki and Perc, 2014; Hilbe et al., 2015b), adaptive dynamics (Hilbe et al., 2013b), and human behavioral experiments (Hilbe et al., 2014a). In each of these studies, the game is assumed to have only two actions: cooperate and defect. In fact, the qualifier “zero-determinant” actually reflects this assumption because these strategies force a matrix determinant to vanish for action spaces with only two options. We show here that this assumption is unnecessary.

More specifically, suppose that players X and Y interact repeatedly with no limit on the number of interactions. For games with two actions, C and D , a memory-one strategy for player X is a vector, $\mathbf{p} = (p_{CC}, p_{CD}, p_{DC}, p_{DD})^\top$, where p_{xy} is the probability that X cooperates following an outcome in which X plays x and Y plays y . Let $\mathbf{s}_X = (R, S, T, P)^\top$ and $\mathbf{s}_Y = (R, T, S, P)^\top$ be the payoff vectors for players X and Y , respectively, and let α , β , and γ be fixed constants. Press and Dyson (2012) show that if there is a constant, ϕ , for which

$$\tilde{\mathbf{p}} := \begin{pmatrix} p_{CC} - 1 \\ p_{CD} - 1 \\ p_{DC} \\ p_{DD} \end{pmatrix} = \phi (\alpha \mathbf{s}_X + \beta \mathbf{s}_Y + \gamma), \quad (2)$$

then X can unilaterally enforce the linear relationship $\alpha\pi_X + \beta\pi_Y + \gamma = 0$ on the average payoffs, π_X and π_Y , by playing \mathbf{p} . A strategy, \mathbf{p} , that satisfies Eq. (2) is known as a “zero-determinant” strategy due to the fact that $\tilde{\mathbf{p}}$ causes a particular matrix determinant to vanish (Press and Dyson, 2012). However, what is important about these strategies is not that they cause some matrix determinant to vanish, but rather that they unilaterally enforce a linear relationship on expected payoffs. Therefore, we refer to these strategies and their generalization to arbitrary action spaces as autocratic strategies. Of particular interest are extortionate strategies, which ensure that a player receives an unfair share of the payoffs exceeding the payoff at the Nash equilibrium (Stewart and Plotkin, 2012). Hence, if P is the payoff for mutual defection in the Prisoner’s Dilemma, then \mathbf{p} is an extortionate strategy for player X if \mathbf{p} enforces the equation $\pi_X - P = \chi (\pi_Y - P)$ for some extortion factor, $\chi \geq 1$.

The most common extensions of finite action sets are continuous action spaces. An element $s \in [0, K]$ represents a player’s investment or cooperation level (up to some maximum, K), such as the amount a player invests in a public good (Doebeli and Hauert, 2005); the volume of blood one vampire bat donates to another (Wilkinson, 1984); the amount of resources used by microbes to produce siderophores (West and Buckling, 2003); or the effort expended in intraspecies grooming (Hemelrijk, 1994; Akinyi et al., 2013). It is important to note that games with continuous action spaces can yield qualitatively different results than their discrete counterparts. For example, the strategy “raise-the-stakes” initially offers a small investment in Prisoner’s Dilemma interactions and subsequently raises the contribution in discrete increments if the opponent matches or betters the investment (Roberts and Sherratt, 1998). However, in a continuous action space, raise-the-stakes evolves into defection due to the fact that another strategy can be arbitrarily close—in terms of the initial investment and subsequent increases in contribution—yet exhibit qualitatively different behavior (Killingback and Doebeli, 1999). In particular, raise-the-stakes succeeds in a discrete action space but fails in a continuous one.

Akin (2015) calls the vector, $\tilde{\mathbf{p}}$, of Eq. (2) a Press-Dyson vector. For continuous action spaces, the payoff vectors, \mathbf{s}_X and \mathbf{s}_Y , must be replaced by payoff functions, $u_X(x, y)$ and $u_Y(x, y)$. That is, $u_i(x, y)$ denotes the payoff to player i when X plays x and Y plays y . The analogue of the linear combination $\alpha \mathbf{s}_X + \beta \mathbf{s}_Y + \gamma$ is the function $\alpha u_X + \beta u_Y + \gamma$. Here, we formally define a Press-Dyson function that extends the Press-Dyson vector to iterated games with arbitrary action spaces. This extension allows one to deduce the existence of strategies that unilaterally enforce linear relationships on the payoffs in more general iterated games. In particular, autocratic (or zero-determinant strategies) are not peculiar to games with two (or even finitely many) actions. Moreover, under mild conditions, player X can enforce a linear relationship on expected payoffs by choosing a memory-one strategy that plays just two actions, despite the fact that the opponent may have an infinite number of actions from which to choose. We give examples of such autocratic strategies in the continuous Donation Game, which represents an instance of the Prisoner’s Dilemma but with an extended, continuous action space.

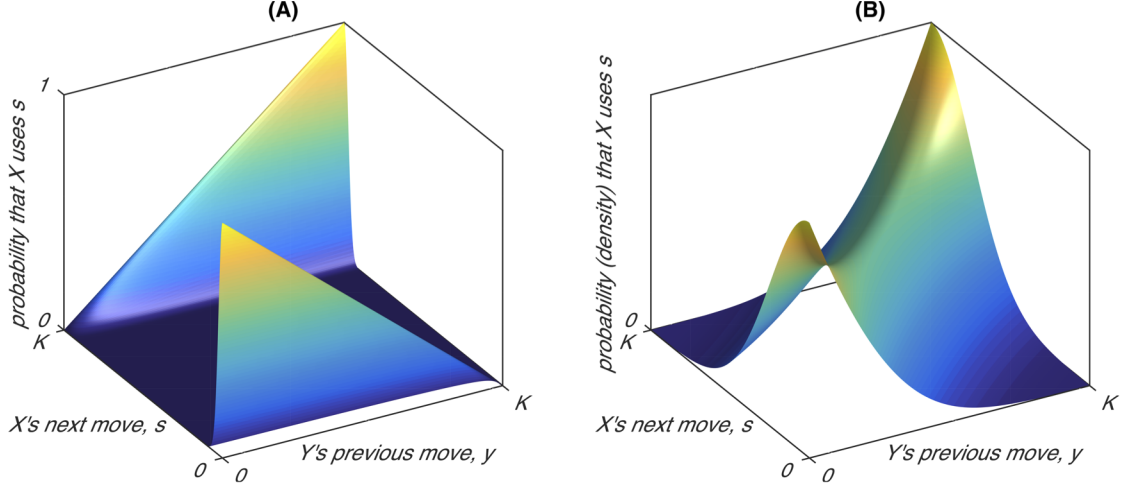


FIGURE 1. Reactive, memory-one strategies, $\sigma_X[x, y](s) = \sigma_X[y](s)$, for a game whose action space is the interval $[0, K]$. In (A), player X uses Y 's action in the previous round to determine the probabilities with which she plays 0 and K in the next round. As Y 's previous action (investment level), y , increases, so does the probability that X uses K in the subsequent round. Since X plays only these two actions, this strategy is called a two-point strategy. In (B), player X uses Y 's action in the previous round to determine the probability density function she uses to devise her next action. In contrast to (A), which depicts a strategy concentrated on just two actions, the strategy depicted in (B) is concentrated on a continuous range of actions in $[0, K]$. As Y 's previous action, y , increases, the mean of the density function governing X 's next action increases, which indicates that X is more willing reciprocate by increasing her investment (action) in response to Y increasing his investment.

2. AUTOCRATIC STRATEGIES

Consider a two-player iterated game with actions spaces, S_X and S_Y , and payoff functions, $u_X(x, y)$ and $u_Y(x, y)$, for players X and Y , respectively. Players X and Y interact repeatedly (infinitely many times), deriving a payoff at each round based on u_X and u_Y . We treat games with temporally-discounted payoffs, which means that for some discounting factor λ with $0 < \lambda < 1$, a payoff of 1 at time $t + \tau$ is treated the same as a payoff of λ^τ at time t (see Fudenberg and Tirole, 1991). Alternatively, one may interpret this game as having a finite number of rounds, where in any given round λ denotes the probability of another round (Nowak, 2006b), which results in an expected game length of $1/(1 - \lambda)$ rounds.

Suppose that, for each $t \geq 0$, x_t and y_t are the actions used by players X and Y at time t . Then, irrespective of the interpretation of λ , the average payoff to player X is

$$\pi_X = (1 - \lambda) \sum_{t=0}^{\infty} \lambda^t u_X(x_t, y_t). \quad (3)$$

The payoff to player Y , π_Y , is obtained by replacing u_X with u_Y in Eq. (3). If the strategies of X and Y are stochastic, then the payoffs are random variables with expectations, π_X and π_Y (see **Supporting Information**). Of particular interest are memory-one strategies, which are probabilistic strategies that depend on only the most recent outcome of the game. If σ_X is a memory-one strategy for X , then we denote by $\sigma_X[x, y](s)$ the probability that X uses s after X plays x and Y plays y (see Fig. 1).

The proofs of the existence of zero-determinant strategies (both in games with and without discounting) rely heavily on the fact that the action space is finite (Press and Dyson, 2012; Hilbe et al., 2015a; Akin, 2015). In particular, it remains unclear whether zero-determinant strategies are consequences of the finiteness of the action space or instances of a more general concept. Here, we introduce autocratic strategies as an

extension of zero-determinant strategies to discounted games with arbitrary action spaces. The traditional, undiscounted case is recovered in the limit $\lambda \rightarrow 1$.

Theorem 1 (Autocratic strategies). Suppose that $\sigma_X[x, y]$ is a memory-one strategy for player X and let σ_X^0 be player X 's initial action. If, for some bounded function, ψ , the equation

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \quad (4)$$

holds for each $x \in S_X$ and $y \in S_Y$, then σ_X^0 and $\sigma_X[x, y]$ together enforce the linear payoff relationship

$$\alpha \pi_X + \beta \pi_Y + \gamma = 0 \quad (5)$$

for *any* strategy of player Y . In other words, the pair $(\sigma_X^0, \sigma_X[x, y])$ is an autocratic strategy for player X .

Note that the initial action, σ_X^0 , in Eq. (4) becomes irrelevant without discounting ($\lambda \rightarrow 1$). The function ψ may be interpreted as a “scaling function” that is used to ensure $\sigma_X[x, y]$ is a feasible memory-one strategy; that is, ψ plays the same role as the scalar ϕ in Eq. (2), which is chosen so that the entries of \mathbf{p} are all between 0 and 1. We call the right-hand side of Eq. (4) a Press-Dyson function, which extends the Press-Dyson vector of Eq. (2) to arbitrary action spaces (see **Supporting Information**). In contrast to action spaces with two options (“cooperate” and “defect”, for instance), autocratic strategies are defined only implicitly via Eq. (4) for general action spaces (and actually already for games with just three actions).

For each x and y , the integral, $\int_{s \in S_X} \psi(s) d\sigma_X[x, y](s)$, may be thought of as the weighted average (expectation) of ψ with respect to $\sigma_X[x, y]$. Since the integral is taken against $\sigma_X[x, y]$, in general one cannot solve Eq. (4) explicitly for $\sigma_X[x, y]$, so it is typically not possible to directly specify all pairs $(\sigma_X^0, \sigma_X[x, y])$ that unilaterally enforce Eq. (5).

Interestingly, under mild conditions, σ_X can be chosen to be a remarkably simple “two-point” strategy, concentrated on just two actions, s_1 and s_2 (see Corollary 2 in **Supporting Information**). Player X can enforce Eq. (5) by playing either s_1 or s_2 in each round, with probabilities determined by the outcome of the previous round (see Fig. 1(A)). Thus, a strategy of this form uses the (memory-one) history of previous play only to adjust the relative weights placed on s_1 and s_2 , while s_1 and s_2 themselves remain unchanged. Unlike in the case of arbitrary σ_X , for fixed ψ , s_1 , and s_2 , it is possible to explicitly solve for all autocratic, two-point strategies on s_1 and s_2 satisfying Eq. (4) (see Remark 3 in **Supporting Information**). In a two-action game, every memory-one strategy is concentrated on two points, which explains why games like the classical Prisoner’s Dilemma fail to capture the implicit nature of autocratic strategies.

3. CONTINUOUS DONATION GAME

In the classical Donation Game, cooperators pay a cost, c , to provide a benefit, b , to the opponent (Sigmund, 2010). Defectors make no donations and pay no costs. The payoff matrix for this game is given by Eq. (1) with $R = b - c$, $S = -c$, $T = b$, and $P = 0$. For $b > c > 0$ a social dilemma arises because the payoff for mutual defection (the Nash equilibrium) is strictly less than the payoff for mutual cooperation, yet both players are tempted to shirk donations, which represents an instance of the Prisoner’s Dilemma. In the iterated (and undiscounted) version of the Donation Game, the main result of Press and Dyson (2012) implies that a memory-one strategy for player X , $\mathbf{p} = (p_{CC}, p_{CD}, p_{DC}, p_{DD})^\top$, enforces $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ for some κ and $\chi \geq 1$ whenever there exists a scalar, ϕ , for which

$$p_{CC} = 1 - \phi(\chi - 1)(b - c - \kappa); \quad (6a)$$

$$p_{CD} = 1 - \phi(\chi b + c - (\chi - 1)\kappa); \quad (6b)$$

$$p_{DC} = \phi(b + \chi c + (\chi - 1)\kappa); \quad (6c)$$

$$p_{DD} = \phi(\chi - 1)\kappa. \quad (6d)$$

The term χ denotes the extortion factor and κ the baseline payoff, i.e. the payoff of \mathbf{p} against itself (Hilbe et al., 2014a). For example, if $\kappa = 0$ and $\chi \geq 1$, then Eq. (6) defines an extortionate strategy, which unilaterally enforces $\pi_X = \chi \pi_Y$ as long as ϕ is sufficiently small. In this sense, ϕ acts as a scaling factor to ensure each coordinate of \mathbf{p} falls between 0 and 1.

Instead of discrete “levels” of cooperation, the continuous Donation Game admits a range of cooperation levels, $[0, K]$, with $K > 0$ indicating maximal cooperation. The costs and benefits associated with s , denoted by $c(s)$ and $b(s)$, respectively, are nondecreasing functions of s and, in analogy to the discrete case, satisfy $b(s) > c(s)$ for $s > 0$ and $c(0) = b(0) = 0$ (Killingback et al., 1999; Wahl and Nowak, 1999a,b; Killingback and Doebeli, 2002; Doebeli et al., 2004). The payoff matrix, Eq. (1) is replaced by payoff functions, with the payoffs to players X and Y for playing x against y being $u_X(x, y) := b(y) - c(x)$ and $u_Y(x, y) = u_X(y, x) = b(x) - c(y)$, respectively (i.e. the game is symmetric). For this natural extension of the classical Donation Game, we first show the existence of autocratic and, in particular, extortionate strategies, that play only $x = 0$ and $x = K$ and ignore all other cooperation levels.

3.1. Two-point autocratic strategies. For the continuous Donation Game, we show, using Theorem 1, that player X can unilaterally enforce $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ for fixed χ and κ by playing only two actions: $x = 0$ (defect) and $x = K$ (fully cooperate). Conditioned on the fact that X plays only 0 and K , a memory-one strategy for player X is defined by a reaction function, $p(x, y)$, which denotes the probability that X plays K following an outcome in which X plays $x \in \{0, K\}$ and Y plays $y \in [0, K]$; $1 - p(x, y)$ is the probability that X plays 0 (i.e. defects). Player X ’s initial action is determined by the probability, p_0 , that X plays $x = K$ in the first round.

Consider the function $\psi(s) := -\chi b(s) - c(s)$ and suppose that $0 \leq \kappa \leq b(K) - c(K)$. If player X ’s initial action is $x = K$ with probability p_0 and $x = 0$ with probability $1 - p_0$ then, for sufficiently weak discounting or, equivalently, sufficiently many rounds of interaction,

$$\lambda \geq \frac{b(K) + \chi c(K)}{\chi b(K) + c(K)}, \quad (7)$$

and for p_0 falling within a feasible range (see Eq. (SI.60) in **Supporting Information**), the reaction function

$$p(x, y) = \frac{1}{\lambda} \left(\frac{b(y) + \chi c(y) + (\chi - 1)\kappa}{\chi b(K) + c(K)} - (1 - \lambda)p_0 \right), \quad (8)$$

defines a memory-one strategy,

$$\sigma_X^0 = (1 - p_0)\delta_0 + p_0\delta_K; \quad (9a)$$

$$\sigma_X[x, y] = (1 - p(x, y))\delta_0 + p(x, y)\delta_K, \quad (9b)$$

where δ_s denotes the Dirac measure centered at $s \in [0, K]$, that enforces the equation $\pi_X - \kappa = \chi(\pi_Y - \kappa)$. If there is no discounting (i.e. $\lambda = 1$), then the initial move is irrelevant and p_0 can be anything in the interval $[0, 1]$. Note that Eq. (8) represents a reactive strategy (Nowak and Sigmund, 1990) because X conditions her play on only the previous move of the opponent (see Fig. 4(B)).

For $\kappa = 0$ and $\chi \geq 1$, Eq. (8) defines an extortionate strategy, σ_X , which guarantees player X a fixed share of the payoffs over the payoff for mutual defection. If $\chi = 1$ (and κ is arbitrary), then this strategy is fair since player X ensures the opponent has a payoff equal to her own (Hilbe et al., 2014b). On the other hand, if $\kappa = b(K) - c(K)$ and $\chi \geq 1$, then Eq. (8) defines a generous (or “compliant”) strategy (Stewart and Plotkin, 2012; Hilbe et al., 2013a). By playing a generous strategy, player X ensures that her payoff is at most that of her opponent’s. For each of these types of strategies, the probability that X reacts to y by cooperating is increasing as a function of y . In particular, X is most likely to cooperate after Y fully cooperates ($y = K$) and is most likely to defect after Y defects ($y = 0$). Moreover, this single choice of $\psi(s) = -\chi b(s) - c(s)$ demonstrates the existence of each of these three classes of autocratic strategies for the continuous Donation Game provided λ is sufficiently weak.

Similarly, if $\psi(s) = b(s)$ and $\lambda \geq c(K)/b(K)$, then

$$p(x, y) = \frac{1}{\lambda} \left(\frac{c(y) + \gamma}{b(K)} - (1 - \lambda)p_0 \right), \quad (10)$$

defines a reaction probability that allows X to set $\pi_Y = \gamma$ for any γ satisfying $0 \leq \gamma \leq b(K) - c(K)$, provided p_0 falls within a suitable range (see Eq. (SI.63) in **Supporting Information**). A strategy that allows a player to single-handedly set the score of the opponent is termed an equalizer strategy (Boerlijst et al., 1997; Hilbe et al., 2013a). However, no autocratic strategy allows player X to set her own score via Eq. (4) (see **Supporting Information**). These results are consistent with the observations of Press and Dyson (2012)




FIGURE 2. Two-point extortionate, generous, and equalizer strategies for the continuous Donation Game with action spaces $S_X = \{0, 2\}$ and $S_Y = [0, 2]$ and discounting factor $\lambda = 0.95$. The costs are linear, $c(s) = 2s$, and the benefits saturating, $b(c(s)) = 5(1 - e^{-c(s)})$. The light blue regions represent the feasible payoff pairs, (π_Y, π_X) , for the repeated, continuous Donation Game when X takes advantage of the entire action space, $[0, 2]$, and the hatched regions represent the feasible payoff pairs when X uses just $\{0, 2\}$. In each panel, the simulation results were obtained by plotting the average payoffs for a fixed autocratic strategy against 1000 randomly chosen memory-one strategies. (A) depicts extortionate strategies enforcing $\pi_X = \chi\pi_Y$ with $\chi = 2$ (black) and $\chi = 3$ (blue). (B) illustrates generous strategies enforcing $b(2) - c(2) - \pi_X = \chi(b(2) - c(2) - \pi_Y)$ with $\chi = 2$ (black) and $\chi = 3$ (blue). (C) demonstrates equalizer strategies enforcing $\pi_Y = \gamma$ with $\gamma = 0$ (black) and $\gamma = b(2) - c(2)$ (blue). From (A) and (B), it is clear that neither mutual extortion nor mutual generosity is a Nash equilibrium since, when both players use extortionate or generous strategies, either player can single-handedly improve the payoffs of both players by deviating from his or her strategy. When both players together use the same action, the best outcome occurs at an investment level of $m = \frac{1}{2} \log 5 < 2$, which results in a payoff of $b(m) - c(m) > b(2) - c(2)$ to both players (yellow dot).

that, in the classical Prisoner's Dilemma without discounting, player X cannot set her own score but can set player Y 's score to anything between the payoffs for mutual defection and mutual cooperation.

3.2. Deterministic autocratic strategies. One feature of two-point autocratic strategies is that they allow a player to exert control over the payoffs of a repeated game while ignoring most of the action space. One drawback is that they restrict the region of feasible game payoffs (see Fig. 2). This shortcoming of two-point strategies leads to a new class of strategies called deterministic strategies, which are perhaps the simplest alternatives to two-point strategies.

A deterministic strategy requires a player to respond to a history of previous play by playing an action with certainty rather than probabilistically. For example, a memory-one deterministic strategy for player X is defined by (i) an initial action, $x_0 \in S_X$, and (ii) a reaction function, $r^X : S_X \times S_Y \rightarrow S_X$, such that X plays $r^X(x, y)$ following an outcome in which X plays x and Y plays y . One well-known example of a deterministic strategy is tit-for-tat in the classical Prisoner's Dilemma, which is defined by $x_0 = C$ (initially cooperate) and $r^X(x, y) = y$ (do what the opponent did in the previous round). Tit-for-tat is also an autocratic strategy since it enforces the fair relationship $\pi_X = \pi_Y$ (Stewart and Plotkin, 2013).

For general memory-one deterministic strategies, the condition for the existence of autocratic strategies, Eq. (4), becomes

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \psi(r^X(x, y)) - (1 - \lambda) \psi(x_0). \quad (11)$$

If Eq. (11) holds for each $x \in S_X$ and $y \in S_Y$, then the deterministic strategy defined by x_0 and r^X enforces $\alpha\pi_X + \beta\pi_Y + \gamma = 0$. Thus, in order to enforce $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ for $0 \leq \kappa \leq b(K) - c(K)$, one may

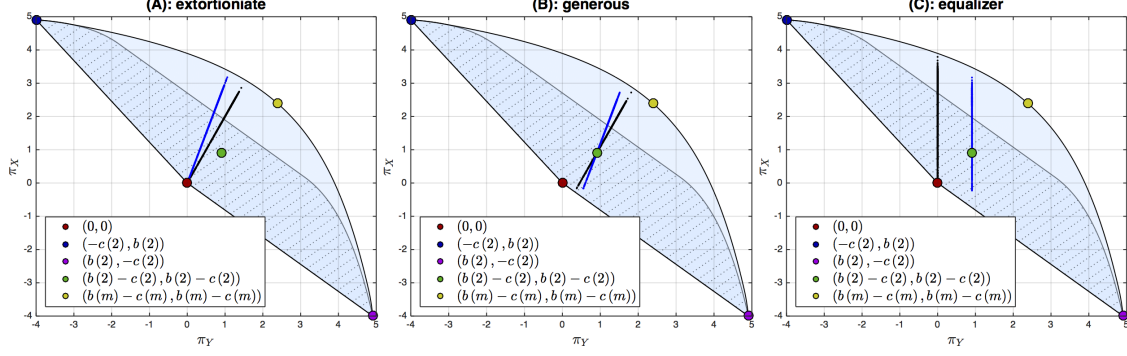


FIGURE 3. Deterministic extortionate, generous, and equalizer strategies for the continuous Donation Game with action spaces $S_X = S_Y = [0, 2]$ and discounting factor $\lambda = 0.95$. As in Fig. 2, the simulation results in each panel were obtained by plotting the average payoffs for a fixed autocratic strategy against 1000 randomly chosen memory-one strategies. (A) and (B) demonstrate extortionate and generous strategies, respectively, with $\chi = 2$ (black) and $\chi = 3$ (blue). (C) shows equalizer strategies enforcing $\pi_Y = \gamma$ with $\gamma = 0$ (black) and $\gamma = b(2) - c(2)$ (blue). Since deterministic strategies allow player X to use a much larger portion of the action space than just 0 and K , we observe that the linear relationships enforced by deterministic autocratic strategies cover a greater portion of the feasible region than do two-point strategies. The best outcome when both players use the same action occurs at $m = \frac{1}{2} \log 5 < 2$, which results in a mutual payoff of $b(m) - c(m) > b(2) - c(2)$ (yellow dot).

choose $\psi(s) := -\chi b(s) - c(s)$ and use the reaction function

$$r^X(x, y) = \psi^{-1} \left(\frac{-b(y) - \chi c(y) - (\chi - 1)\kappa - (1 - \lambda)\psi(x_0)}{\lambda} \right), \quad (12)$$

where $\psi^{-1}(\dots)$ denotes the inverse of the function ψ , provided λ satisfies Eq. (7) and x_0 falls within a feasible range (see Eq. (SI.65) in **Supporting Information**).

Similarly, to enforce $\pi_Y = \gamma$ for $0 \leq \gamma \leq b(K) - c(K)$, one may choose $\psi(s) := b(s)$ and use the reaction function

$$r^X(x, y) = \psi^{-1} \left(\frac{c(y) + \gamma - (1 - \lambda)\psi(x_0)}{\lambda} \right), \quad (13)$$

provided $\lambda \geq c(K)/b(K)$ and, again, x_0 falls within a feasible range (see Eq. (SI.67) in **Supporting Information**).

Examples of deterministic extortionate, generous, and equalizer strategies are given in Fig. 3. It is evident that deterministic strategies increase the feasible region in which linear payoff relationships can be enforced as compared to their two-point counterparts (c.f. Fig. 2).

4. DISCUSSION

In games with two actions, zero-determinant strategies are typically defined via a technical condition such as Eq. (2) (Press and Dyson, 2012; Hilbe et al., 2015a). This definition makes generalizations to games with larger action spaces difficult because Eq. (2) makes sense only for two-action games. Therefore, we introduce the more general term autocratic strategy for any strategy that unilaterally enforces a linear relationship on expected payoffs. Of course, this linear relationship is precisely what makes strategies satisfying Eq. (2) interesting.

Theorem 1 provides a condition for the existence of autocratic strategies for games with general action spaces. We illustrate this phenomenon with a continuous-action-space extension of the classical Donation Game, which represents an instance of the Prisoner's Dilemma. The existing literature on zero-determinant strategies for the classical Prisoner's Dilemma provides no way of treating this continuous extension of the

Donation Game. However, Theorem 1 makes no assumptions on the action space of the game and thus applies to the continuous Donation Game as well as its classical counterpart.

Surprisingly, in many cases a player can enforce a linear relationship on expected payoffs by playing only two actions, despite the fact that the opponent may have infinitely many actions available to use (Corollary 2 in **Supporting Information**). We demonstrate that the conditions guaranteeing the existence of extortionate, generous, fair, and equalizer strategies in the continuous Donation Game are in fact similar to those of the two-action case. However, despite the simplicity of these two-point strategies, a player needs to know how to respond to every possible move of the opponent; knowledge of how to respond to just defection ($y = 0$) and full cooperation ($y = K$) does not suffice. Therefore, although a player using a two-point strategy plays only $x = 0$ and $x = K$, these strategies represent a departure from the classical Donation Game.

Another important difference is that, whereas in the classical Prisoner’s Dilemma mutual generosity represents a symmetric Nash equilibrium (Hilbe et al., 2015a), this need not be the case in the continuous Donation Game. Instead, mutual generosity results in a payoff of $b(K) - c(K)$ for each player, where K is the maximal investment, but intermediate levels of cooperation may yield $b(m) - c(m) > b(K) - c(K)$ for $m \in (0, K)$, i.e. both players fare better if they each invest m instead of K (see Figs. 2 and 3). However, no player can enforce a generous relationship with baseline payoff $\kappa = b(m) - c(m)$ because it is outside of the feasible range for κ (see **Supporting Information**). Thus, the performance of a generous strategy as a response to itself depends critically on whether the game has two actions or a continuous range of actions.

Extortionate strategies for the iterated Prisoner’s Dilemma are not evolutionarily stable (Adami and Hintze, 2013). Since mutual generosity in the continuous Donation Game need not be a Nash equilibrium, it follows that generous strategies also need not be evolutionarily stable. Moreover, against human opponents, extortioners are punished by a refusal to fully cooperate, while generous players provide their opponents with an incentive to cooperate and fare better in experiments (Hilbe et al., 2014a). Such behavior supports what one would expect from a player using a fair autocratic strategy enforcing $\pi_X = \pi_Y$ (such as tit-for-tat): if the opponent ensures $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ for some $\chi > 1$, then both players get κ . In particular, fair strategies punish extortion and reward generosity. Based on our results on generous strategies for the continuous Donation Game, it would be interesting to see whether (human) experiments support the same conclusion and whether the participants succeed in securing payoffs that exceed those of mutual generosity. The performance of autocratic strategies in populations, however, is but one perspective on this recently-discovered class of strategies for repeated games.

In games with two discrete actions, our definition of a Press-Dyson function specializes to a multiple of the Press-Dyson vector, $C\tilde{\mathbf{p}}$ with C a constant (see Eq. (2) and **Supporting Information** for details). The Press-Dyson vector is recovered by normalizing the Press-Dyson function and thus eliminating the constant, C . However, in games with d actions, this function involves at least $d - 1$ constants, which, for $d > 2$, cannot be eliminated by normalization (see **Supporting Information** for an example with three actions). Therefore, based on Theorem 1, in two-action games it is perhaps more appropriate to define a Press-Dyson vector to be any vector of the form $C\tilde{\mathbf{p}}$. This distinction for games with two actions is minor, however, and does not change the fact that Theorem 1 covers all of the known results on the existence of zero-determinant strategies for repeated games.

More importantly, however, the analysis of iterated games with only two actions completely misses the fact that autocratic strategies are most naturally presented implicitly via Eq. (4). Even in the case of two actions, infinitely many autocratic strategies may exist (Press and Dyson, 2012), but their simplistic nature admits explicit solutions. Our extension shows that, in general, (i) autocratic strategies need not be unique and (ii) one cannot explicitly list all autocratic strategies that produce a fixed Press-Dyson function. Thus, for arbitrary action spaces (but already for games with $d > 2$ actions), the space of autocratic strategies is more sophisticated than two-action games suggest. Notwithstanding the intrinsic difficulty in explicitly specifying all autocratic strategies, our results demonstrate that these strategies exist in a broad class of games and are not simply consequences of the finiteness of the action space in games such as the Prisoner’s Dilemma.

ACKNOWLEDGMENTS

The authors are grateful to Christian Hilbe for reading an earlier draft of this manuscript and for suggesting that we look into deterministic strategies. The authors acknowledge financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC), grant RGPIN-2015-05795.

REFERENCES

- C. Adami and A. Hintze. Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything. *Nature Communications*, 4, Aug 2013. doi: 10.1038/ncomms3193.
- E. Akin. What You Gotta Know to Play Good in the Iterated Prisoner’s Dilemma. *Games*, 6(3):175–190, Jun 2015. doi: 10.3390/g6030175.
- M. Y. Akinyi, J. Tung, M. Jeneby, N. B. Patel, J. Altmann, and S. C. Alberts. Role of grooming in reducing tick load in wild baboons (*Papio cynocephalus*). *Animal Behaviour*, 85(3):559–568, Mar 2013. doi: 10.1016/j.anbehav.2012.12.012.
- R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, Mar 1981. doi: 10.1126/science.7466396.
- M. C. Boerlijst, M. A. Nowak, and K. Sigmund. Equal pay for all prisoners. *The American Mathematical Monthly*, 104(4):303, Apr 1997. doi: 10.2307/2974578.
- P. Dal Bó. Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games. *American Economic Review*, 95(5):1591–1604, Dec 2005. doi: 10.1257/000282805775014434.
- R. M. Dawes. Social dilemmas. *Annual Review of Psychology*, 31(1):169–193, Jan 1980. doi: 10.1146/annurev.ps.31.020180.001125.
- A. W. Delton, M. M. Krasnow, L. Cosmides, and J. Tooby. Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proceedings of the National Academy of Sciences*, 108(32):13335–13340, Jul 2011. doi: 10.1073/pnas.1102131108.
- M. Doebeli and C. Hauert. Models of cooperation based on the prisoner’s dilemma and the snowdrift game. *Ecology Letters*, 8(7):748–766, Jul 2005. doi: 10.1111/j.1461-0248.2005.00773.x.
- M. Doebeli, C. Hauert, and T. Killingback. The Evolutionary Origin of Cooperators and Defectors. *Science*, 306(5697):859–862, Oct 2004. doi: 10.1126/science.1101456.
- D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.
- J. B. Heide and A. S. Miner. The shadow of the future : Effects of anticipated interaction and frequency of contact on buy-seller cooperation. *Academy of Management Journal*, 35(2):265–291, Jun 1992. doi: 10.2307/256374.
- C. K. Hemelrijk. Support for being groomed in long-tailed macaques, *Macaca fascicularis*. *Animal Behaviour*, 48(2):479–481, Aug 1994. doi: 10.1006/anbe.1994.1264.
- C. Hilbe, M. A. Nowak, and K. Sigmund. Evolution of extortion in Iterated Prisoner’s Dilemma games. *Proceedings of the National Academy of Sciences*, 110:6913–6918, Apr 2013a. doi: 10.1073/pnas.1214834110.
- C. Hilbe, M. A. Nowak, and A. Traulsen. Adaptive Dynamics of Extortion and Compliance. *PLoS ONE*, 8(11):e77886, Nov 2013b. doi: 10.1371/journal.pone.0077886.
- C. Hilbe, T. Röhl, and M. Milinski. Extortion subdues human players but is finally punished in the prisoner’s dilemma. *Nature Communications*, 5, May 2014a. doi: 10.1038/ncomms4976.
- C. Hilbe, B. Wu, A. Traulsen, and M. A. Nowak. Cooperation and control in multiplayer social dilemmas. *Proceedings of the National Academy of Sciences*, 111(46):16425–16430, Oct 2014b. doi: 10.1073/pnas.1407887111.
- C. Hilbe, A. Traulsen, and K. Sigmund. Partners or rivals? Strategies for the iterated prisoner’s dilemma. *Games and Economic Behavior*, 92:41–52, Jul 2015a. doi: 10.1016/j.geb.2015.05.005.
- C. Hilbe, B. Wu, A. Traulsen, and M. A. Nowak. Evolutionary performance of zero-determinant strategies in multiplayer games. *Journal of Theoretical Biology*, 374:115–124, Jun 2015b. doi: 10.1016/j.jtbi.2015.03.032.
- T. Killingback and M. Doebeli. ‘Raise the stakes’ evolves into a defector. *Nature*, 400(6744):518–518, Aug 1999. doi: 10.1038/22913.

- T. Killingback and M. Doebeli. The Continuous Prisoner's Dilemma and the Evolution of Cooperation through Reciprocal Altruism with Variable Investment. *The American Naturalist*, 160(4):421–438, Oct 2002. doi: 10.1086/342070.
- T. Killingback, M. Doebeli, and N. Knowlton. Variable investment, the Continuous Prisoner's Dilemma, and the origin of cooperation. *Proceedings of the Royal Society B: Biological Sciences*, 266(1430):1723–1728, Sep 1999. doi: 10.1098/rspb.1999.0838.
- J. Korevaar. *Tauberian Theory*. Springer Berlin Heidelberg, 2004. doi: 10.1007/978-3-662-10225-1.
- M. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, 20(3):247–265, Sep 1990. doi: 10.1007/bf00049570.
- M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364(6432):56–58, Jul 1993. doi: 10.1038/364056a0.
- M. A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006a.
- M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, Dec 2006b. doi: 10.1126/science.1133755.
- L. Pan, D. Hao, Z. Rong, and T. Zhou. Zero-Determinant Strategies in Iterated Public Goods Game. *Scientific Reports*, 5:13096, Aug 2015. doi: 10.1038/srep13096.
- W. H. Press and F. J. Dyson. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences*, 109(26):10409–10413, May 2012. doi: 10.1073/pnas.1206569109.
- G. Roberts and T. N. Sherratt. Development of cooperative relationships through increasing investment. *Nature*, 394(6689):175–179, Jul 1998. doi: 10.1038/28160.
- K. Sigmund. *The calculus of selfishness*. Princeton University Press, 2010.
- A. J. Stewart and J. B. Plotkin. Extortion and cooperation in the Prisoner's Dilemma. *Proceedings of the National Academy of Sciences*, 109(26):10134–10135, Jun 2012. doi: 10.1073/pnas.1208087109.
- A. J. Stewart and J. B. Plotkin. From extortion to generosity, evolution in the Iterated Prisoner's Dilemma. *Proceedings of the National Academy of Sciences*, 110(38):15348–15353, Sep 2013. doi: 10.1073/pnas.1306246110.
- A. Szolnoki and M. Perc. Evolution of extortion in structured populations. *Physical Review E*, 89(2), Feb 2014. doi: 10.1103/physreve.89.022804.
- R. L. Trivers. The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(1):35–57, Mar 1971. doi: 10.1086/406755.
- P. A. M. Van Lange, J. Joireman, C. D. Parks, and E. Van Dijk. The psychology of social dilemmas: A review. *Organizational Behavior and Human Decision Processes*, 120(2):125–141, Mar 2013. doi: 10.1016/j.obhdp.2012.11.003.
- L. M. Wahl and M. A. Nowak. The Continuous Prisoner's Dilemma: I. Linear Reactive Strategies. *Journal of Theoretical Biology*, 200(3):307–321, Oct 1999a. doi: 10.1006/jtbi.1999.0996.
- L. M. Wahl and M. A. Nowak. The Continuous Prisoner's Dilemma: II. Linear Reactive Strategies with Noise. *Journal of Theoretical Biology*, 200(3):323–338, Oct 1999b. doi: 10.1006/jtbi.1999.0997.
- S. A. West and A. Buckling. Cooperation, virulence and siderophore production in bacterial parasites. *Proceedings of the Royal Society B: Biological Sciences*, 270(1510):37–44, Jan 2003. doi: 10.1098/rspb.2002.2209.
- G. S. Wilkinson. Reciprocal food sharing in the vampire bat. *Nature*, 308(5955):181–184, Mar 1984. doi: 10.1038/308181a0.

SUPPORTING INFORMATION

SI.1. ITERATED GAMES WITH TWO PLAYERS AND MEASURABLE ACTION SPACES

By “action space,” we mean a measurable space, S , equipped with a σ -algebra, $\mathcal{F}(S)$ (although we suppress $\mathcal{F}(S)$ and refer to the space simply as S). Informally, S is the space of actions, decisions, investments, or options available to a player at each round of the iterated interaction and could be a finite set, a continuous interval, or something more complicated. Since the players need not have the same action space, we denote by S_X the space of actions available to player X and by S_Y the space of actions available to player Y . In what follows, all functions are assumed to be measurable and bounded.

In each encounter (i.e. “one-shot game”), the players receive payoffs based on a payoff function,

$$u = (u_X, u_Y) : S_X \times S_Y \longrightarrow \mathbb{R}^2. \quad (\text{SI.1})$$

The first and second coordinate functions, u_X and u_Y , give the payoffs to players X and Y , respectively. An iterated game between players X and Y consists of a sequence of these one-shot interactions. If, at time t , player X uses $x_t \in S_X$ and player Y uses $y_t \in S_Y$, then the (normalized) payoff to player X for a sequence of $T + 1$ interactions (from time $t = 0$ to $t = T$) is

$$\frac{1 - \lambda}{1 - \lambda^{T+1}} \sum_{t=0}^T \lambda^t u_X(x_t, y_t), \quad (\text{SI.2})$$

where λ is the discounting factor, $0 < \lambda < 1$. The payoff to player Y is obtained by replacing u_X by u_Y in Eq. (SI.2). Thus, the discounted payoffs, $\lambda^t u_X(x_t, y_t)$, are simply added up and then normalized by a factor of $\frac{1-\lambda}{1-\lambda^{T+1}}$ to ensure that the payoff for the repeated game is measured in the same units as the payoffs for individual encounters (Fudenberg and Tirole, 1991). Moreover, provided the series $\sum_{t=0}^{\infty} u_X(x_t, y_t)$ is Cesàro summable, meaning $\lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T u_X(x_t, y_t)$ exists, we have

$$\lim_{\lambda \rightarrow 1^-} (1 - \lambda) \sum_{t=0}^{\infty} \lambda^t u_X(x_t, y_t) = \lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T u_X(x_t, y_t) \quad (\text{SI.3})$$

(see Korevaar, 2004). Therefore, payoffs in the undiscounted case may be obtained from the payoffs for discounted games in the limit $\lambda \rightarrow 1^-$, provided this limit exists (see Hilbe et al., 2015a).

Here, we consider stochastic strategies that condition on the history of play: both players observe the sequence of play up to the current period and use it to devise an action for the present encounter. In order to formally define such strategies, we first recall the notion of “history” in a repeated game: a history at time T is a sequence of action pairs,

$$h^T := \left((x_0, y_0), \dots, (x_{T-1}, y_{T-1}) \right), \quad (\text{SI.4})$$

indicating the sequence of play leading up to the interaction at time T (Fudenberg and Tirole, 1991). In other words, a history at time T is an element of $\mathcal{H}^T := \prod_{t=0}^{T-1} S_X \times S_Y$. Let $\mathcal{H}^0 := \{\emptyset\}$, where \emptyset denotes the “empty” history (which serves just to indicate that there has been no history of past play, i.e. that the game has not yet begun). The set of all possible histories is the disjoint union, $\mathcal{H} := \bigsqcup_{T \geq 0} \mathcal{H}^T$. For $h^T = \left((x_0, y_0), \dots, (x_{T-1}, y_{T-1}) \right)$ and $t \leq T - 1$, let

$$h_t^T := (x_t, y_t); \quad (\text{SI.5a})$$

$$h_{\leq t}^T := \left((x_0, y_0), \dots, (x_t, y_t) \right). \quad (\text{SI.5b})$$

That is, h_t^T is the action pair played at time t , and $h_{\leq t}^T$ is the “sub-history” of h^T until time $t \leq T$.

A pure strategy for player X in the repeated game is a map, $\mathcal{H} \rightarrow S_X$, indicating an action in S_X (deterministically) for each history leading up to the current encounter. More generally, X could look at the history of past play and use this information to choose an action from S_X probabilistically (rather than deterministically). A strategy of this form is known as a behavioral strategy (Fudenberg and Tirole, 1991). In terms of \mathcal{H} , a behavioral strategy for player X is a map

$$\sigma_X : \mathcal{H} \longrightarrow \Delta(S_X), \quad (\text{SI.6})$$

where $\Delta(S_X)$ is the space of probability measures on S_X . An important type of behavioral strategy is a Markov strategy, which is a behavioral strategy, σ_X , that satisfies $\sigma_X[h^T] = \sigma_X[h_{T-1}^T]$. That is, a Markov strategy depends on only the last pair of actions and not on the entire history of play. Note, however, that a Markov strategy may still depend on t . If σ_X is a Markov strategy that does not depend on t , then we say that σ_X is a stationary (or memory-one) strategy.

Suppose that σ_X and σ_Y are behavioral strategies for players X and Y , respectively. Consider the map, σ , defined by the product measure,

$$\begin{aligned} \sigma &:= \sigma_X \times \sigma_Y : \mathcal{H} \longrightarrow \Delta(S_X \times S_Y) \\ &: h^t \longmapsto \sigma_X[h^t] \times \sigma_Y[h^t]. \end{aligned} \quad (\text{SI.7})$$

By the Hahn-Kolmogorov theorem, for each $t \geq 0$ there exists a unique measure, μ_t , on \mathcal{H}^{t+1} such that for each $E' \in \mathcal{F}(\mathcal{H}^t)$ and $E \in \mathcal{F}(S_X \times S_Y)$,

$$\mu_t(E' \times E) = \int_{h^t \in E'} \sigma(h^t, E) \, d\sigma(h_{\leq t-2}^t, h_{t-1}^t) \cdots d\sigma(h_{\leq 0}^t, h_1^t) \, d\sigma(\emptyset, h_0^t), \quad (\text{SI.8})$$

where, for $h \in \mathcal{H}$ and $s \in S_X \times S_Y$, $d\sigma(h, s)$ denotes the differential of the measure $\sigma(h, -)$ on $S_X \times S_Y$. In the case $t = 0$, this measure is simply the product of the two initial actions, i.e. $\mu_0 = \sigma_X[\emptyset] \times \sigma_Y[\emptyset]$. From these measures, we obtain a sequence of measures, $\{\nu_t\}_{t \geq 0} \subseteq \Delta(S_X \times S_Y)$, defined by

$$\nu_t(E) := \mu_t(\mathcal{H}^t \times E). \quad (\text{SI.9})$$

Informally, $\nu_t(E)$ is the probability that the action pair at time t is in $E \subseteq S_X \times S_Y$, averaged over all histories that lead to E . The sequences, $\{\mu_t\}_{t \geq 0}$ and $\{\nu_t\}_{t \geq 0}$, admit a convenient format for the expected payoffs, π_X^T and π_Y^T , to players X and Y , respectively. Before stating this result, we first formally define expected payoffs for the $(T+1)$ -period game (where $T < \infty$):

Definition 1 (Objective function for a finite game). If σ_X and σ_Y are behavioral strategies for players X and Y , respectively, and if $\sigma = \sigma_X \times \sigma_Y$ (see Eq. (SI.7)), then the objective function (or expected payoff) for player X in the $(T+1)$ -period game is

$$\pi_X^T := \int_{h^{T+1} \in \mathcal{H}^{T+1}} \left[\frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t u_X(h_t^{T+1}) \right] d\sigma(h_{\leq T-1}^{T+1}, h_T^{T+1}) \cdots d\sigma(h_{\leq 0}^{T+1}, h_1^{T+1}) \, d\sigma(\emptyset, h_0^{T+1}). \quad (\text{SI.10})$$

Using $\{\nu_t\}_{t \geq 0}$, we can write π_X^T differently:

Lemma 1. For fixed σ_X and σ_Y generating $\{\nu_t\}_{t \geq 0}$, we have

$$\pi_X^T = \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{(x,y) \in S_X \times S_Y} u_X(x, y) \, d\nu_t(x, y). \quad (\text{SI.11})$$

As a consequence of Lemma 1, we see that $\lim_{T \rightarrow \infty} \pi_X^T$ exists since ν_t is a probability measure and

$$\left| \int_{(x,y) \in S_X \times S_Y} u_X(x, y) \, d\nu_t(x, y) \right| \leq \sup_{(x,y) \in S_X \times S_Y} |u_X(x, y)| < \infty \quad (\text{SI.12})$$

by the fact that u_X is bounded. Thus, we define the objective function for an infinite game as follows:

Definition 2 (Objective function for an infinite game). If σ_X and σ_Y are behavioral strategies for players X and Y , respectively, and if $\sigma = \sigma_X \times \sigma_Y$, then the objective function for player X in the infinite game is

$$\pi_X := \lim_{T \rightarrow \infty} \pi_X^T = (1-\lambda) \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} u_X(x, y) \, d\nu_t(x, y). \quad (\text{SI.13})$$

Remark 1. Classically, the objective function of a repeated game with infinitely many rounds is defined using a distribution over “infinite histories,” which is generated by the players’ strategies for the repeated

game (Fudenberg and Tirole, 1991). That is, for $\mathcal{H}^\infty := \prod_{t=0}^\infty S_X \times S_Y$ and some measure, $\mu \in \Delta(\mathcal{H}^\infty)$, the objective function of player X is defined by

$$\int_{h^\infty \in \mathcal{H}^\infty} (1 - \lambda) \sum_{t=0}^\infty \lambda^t u_X(h_t^\infty) d\mu(h^\infty). \quad (\text{SI.14})$$

Using Eq. (SI.13) as an objective function for player X , we do not need to worry about what μ is (or if it even exists for a general action space) since Eq. (SI.14), whenever it is defined, must coincide with Eq. (SI.13). To see why, suppose that there is a distribution, μ , on \mathcal{H}^∞ that satisfies

$$\mu(E \times (S_X \times S_Y) \times (S_X \times S_Y) \times \cdots) = \mu_T(E) \quad (\text{SI.15})$$

for each $E \in \mathcal{F}(\mathcal{H}^{T+1})$. Then, by the dominated convergence theorem, Eq. (SI.15), and Eq. (SI.9),

$$\begin{aligned} & \int_{h^\infty \in \mathcal{H}^\infty} (1 - \lambda) \sum_{t=0}^\infty \lambda^t u_X(h_t^\infty) d\mu(h^\infty) \\ &= (1 - \lambda) \sum_{t=0}^\infty \lambda^t \int_{h^\infty \in \mathcal{H}^\infty} u_X(h_t^\infty) d\mu(h^\infty) \\ &= (1 - \lambda) \sum_{t=0}^\infty \lambda^t \int_{h^{t+1} \in \mathcal{H}^{t+1}} u_X(h_t^{t+1}) d\mu_t(h^{t+1}) \\ &= (1 - \lambda) \sum_{t=0}^\infty \lambda^t \int_{(x,y) \in S_X \times S_Y} u_X(x, y) d\nu_t(x, y). \end{aligned} \quad (\text{SI.16})$$

Therefore, assuming Lemma 1, the objective function for player X defined by $\pi_X := \lim_{T \rightarrow \infty} \pi_X^T$ is the same as the classical objective function for repeated games when the players' strategies produce a probability distribution over infinite histories. Typically, the existence of such a distribution depends on S being finite or the measures in $\{\mu_t\}_{t \geq 0}$ being inner regular (which allows one to deduce the existence of μ from $\{\mu_t\}_{t \geq 0}$ using the Kolmogorov extension theorem). In practice, these assumptions are often not unreasonable, but with Lemma 1, we do not need to worry about the existence of such a distribution.

In order to prove Lemma 1, we first need a simple technical result:

Lemma 2. Suppose that \mathcal{X} and \mathcal{Y} are measure spaces and σ is a Markov kernel from \mathcal{X} to \mathcal{Y} . Let μ be a probability measure on \mathcal{X} and consider the measure on \mathcal{Y} defined, for $E \in \mathcal{F}(\mathcal{Y})$, by

$$\nu(E) := \int_{x \in \mathcal{X}} \sigma(x, E) d\mu(x). \quad (\text{SI.17})$$

For any bounded, measurable function, $f : \mathcal{Y} \rightarrow \mathbb{R}$, we have

$$\int_{y \in \mathcal{Y}} f(y) d\nu(y) = \int_{x \in \mathcal{X}} \int_{y \in \mathcal{Y}} f(y) d\sigma(x, y) d\mu(x), \quad (\text{SI.18})$$

where, for each $x \in \mathcal{X}$, $d\sigma(x, y)$ denotes the differential of the measure $\sigma(x, -)$ on \mathcal{Y} .

Proof. Since f is bounded, there exists a sequence of simple functions, $\{f_n\}_{n \geq 1}$, such that $f_n \rightarrow f$ uniformly on \mathcal{Y} . For each $n \geq 1$, let $f_n = \sum_{i=1}^{N_n} c_i^n \chi_{E_i^n}$ for some $c_i^n \in \mathbb{R}$ and $E_i^n \in \mathcal{F}(\mathcal{Y})$, where $\chi_{E_i^n}$ is the characteristic

function of E_i^n (meaning $\chi_{E_i^n}(x) = 1$ if $x \in E_i^n$ and $\chi_{E_i^n}(x) = 0$ if $x \notin E_i^n$). By uniform convergence,

$$\begin{aligned}
\int_{y \in \mathcal{Y}} f(y) d\nu(y) &= \lim_{n \rightarrow \infty} \int_{y \in \mathcal{Y}} f_n(y) d\nu(y) \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{N_n} c_i^n \nu(E_i^n) \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{N_n} c_i^n \int_{x \in \mathcal{X}} \sigma(x, E_i^n) d\mu(x) \\
&= \lim_{n \rightarrow \infty} \int_{x \in \mathcal{X}} \sum_{i=1}^{N_n} c_i^n \sigma(x, E_i^n) d\mu(x) \\
&= \lim_{n \rightarrow \infty} \int_{x \in \mathcal{X}} \int_{y \in \mathcal{Y}} f_n(y) d\sigma(x, y) d\mu(x) \\
&= \int_{x \in \mathcal{X}} \int_{y \in \mathcal{Y}} f(y) d\sigma(x, y) d\mu(x), \tag{SI.19}
\end{aligned}$$

which completes the proof. \square

Proof of Lemma 1. By Lemma 2 and the definitions of μ_t and ν_t ,

$$\begin{aligned}
\pi_X^T &= \int_{h^{T+1} \in \mathcal{H}^{T+1}} \left[\frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t u_X(h_t^{T+1}) \right] d\sigma(h_{\leq T-1}^{T+1}, h_T^{T+1}) \cdots d\sigma(h_{\leq 0}^{T+1}, h_1^{T+1}) d\sigma(\emptyset, h_0^{T+1}) \\
&= \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{h^{T+1} \in \mathcal{H}^{T+1}} u_X(h_t^{T+1}) d\sigma(h_{\leq T-1}^{T+1}, h_T^{T+1}) \cdots d\sigma(h_{\leq 0}^{T+1}, h_1^{T+1}) d\sigma(\emptyset, h_0^{T+1}) \\
&= \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{h^{T+1} \in \mathcal{H}^{T+1}} u_X(h_t^{T+1}) d\mu_T(h^{T+1}) \\
&= \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{h^{t+1} \in \mathcal{H}^{t+1}} u_X(h_t^{t+1}) d\mu_t(h^{t+1}) \\
&= \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{h^{t+1} \in S_X \times S_Y} u_X(h_t^{t+1}) d\nu_t(h^{t+1}) \\
&= \frac{1-\lambda}{1-\lambda^{T+1}} \sum_{t=0}^T \lambda^t \int_{(x,y) \in S_X \times S_Y} u_X(x, y) d\nu_t(x, y), \tag{SI.20}
\end{aligned}$$

which completes the proof. \square

The objective function of Eq. (SI.13), which is obtained using Lemma 1, eliminates the need to deal with histories when proving our main results for iterated games. With the background on expected payoffs now established, we turn our attention to the proofs of the results claimed in the main text:

SI.2. DETAILED PROOFS OF THE MAIN RESULTS

Before proving our main results, we state a technical lemma that generalizes Lemma 3.1 of Akin (2015)–which Hilbe et al. (2014b) refer to as Akin’s Lemma–and Lemma 1 of Hilbe et al. (2015a). This lemma relates the strategies of the two players, σ_X and σ_Y , and the (discounted) sequence of play to the initial action of player X when σ_X is memory-one. Our proof of this lemma is essentially the same as theirs but in the broader setting of a measurable action space:

Lemma 3. For any memory-one strategy, σ_X , and $E \in \mathcal{F}(S_X)$,

$$\sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\chi_E(x) - \lambda \sigma_X[x, y](E) \right] d\nu_t(x, y) = \sigma_X^0(E), \tag{SI.21}$$

where σ_X^0 is the initial action of player X .

Proof. By the definition of ν_t , we have

$$\int_{(x,y) \in S_X \times S_Y} \chi_E(x) d\nu_t(x,y) = \nu_t(E \times S); \quad (\text{SI.22a})$$

$$\int_{(x,y) \in S_X \times S_Y} \sigma_X[x,y](E) d\nu_t(x,y) = \nu_{t+1}(E \times S). \quad (\text{SI.22b})$$

Therefore, since ν_t is a probability measure (in particular, at most 1 on any measurable set) for each t ,

$$\begin{aligned} & \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} [\chi_E(x) - \lambda \sigma_X[x,y](E)] d\nu_t(x,y) \\ &= \sum_{t=0}^{\infty} \lambda^t (\nu_t(E \times S) - \lambda \nu_{t+1}(E \times S)) \\ &= \nu_0(E \times S) - \lim_{t \rightarrow \infty} \lambda^{t+1} \nu_{t+1}(E \times S) \\ &= \nu_0(E \times S) \\ &= \sigma_X^0(E), \end{aligned} \quad (\text{SI.23})$$

which completes the proof. \square

By the definitions of π_X and π_Y , we have

$$\alpha \pi_X + \beta \pi_Y + \gamma = (1 - \lambda) \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} [\alpha u_X(x,y) + \beta u_Y(x,y) + \gamma] d\nu_t(x,y). \quad (\text{SI.24})$$

Since our goal is to establish Theorem 1, which states that player X can enforce the relation $\alpha \pi_X + \beta \pi_Y + \gamma = 0$ using some $\sigma_X[x,y]$ and σ_X^0 , as a first step we show that $\sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \varphi(x,y) d\nu_t(x,y) = \int_{s \in S_X} \psi(s) d\sigma_X^0(s)$ for a particular choice of $\varphi(x,y)$. We then deduce Theorem 1 by setting

$$\alpha u_X(x,y) + \beta u_Y(x,y) + \gamma + (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) = \varphi(x,y) \quad (\text{SI.25})$$

for this known function, φ .

Proposition 1. If $\psi : S_X \rightarrow \mathbb{R}$ is a bounded, measurable function, then

$$\sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x,y](s) \right] d\nu_t(x,y) = \int_{s \in S_X} \psi(s) d\sigma_X^0(s), \quad (\text{SI.26})$$

for any memory-one strategy, σ_X , where σ_X^0 is the initial action of player X .

Proof. Since ψ is bounded, there exists a sequence of simple functions, $\{\psi_n\}_{n \geq 1}$, such that $\psi_n \rightarrow \psi$ uniformly on S . For each $n \geq 1$, let $\psi_n = \sum_{i=1}^{N_n} c_i^n \chi_{E_i^n}$. Using the uniform convergence of this sequence, together with

the dominated convergence theorem and Lemma 3, we obtain

$$\begin{aligned}
& \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) \\
&= \sum_{t=0}^{\infty} \lambda^t \lim_{n \rightarrow \infty} \int_{(x,y) \in S_X \times S_Y} \left[\psi_n(x) - \lambda \int_{s \in S_X} \psi_n(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) \\
&= \lim_{n \rightarrow \infty} \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\psi_n(x) - \lambda \int_{s \in S_X} \psi_n(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) \\
&= \lim_{n \rightarrow \infty} \sum_{t=0}^{\infty} \lambda^t \sum_{i=1}^{N_n} c_i^n \int_{(x,y) \in S_X \times S_Y} \left[\chi_{E_i^n}(x) - \lambda \int_{s \in S_X} \chi_{E_i^n}(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) \\
&= \lim_{n \rightarrow \infty} \sum_{t=0}^{\infty} \lambda^t \sum_{i=1}^{N_n} c_i^n \int_{(x,y) \in S_X \times S_Y} \left[\chi_{E_i^n}(x) - \lambda \sigma_X[x, y](E_i^n) \right] d\nu_t(x, y) \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{N_n} c_i^n \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\chi_{E_i^n}(x) - \lambda \sigma_X[x, y](E_i^n) \right] d\nu_t(x, y) \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{N_n} c_i^n \sigma_X^0(E_i^n) \\
&= \lim_{n \rightarrow \infty} \int_{s \in S_X} \psi_n(s) d\sigma_X^0(s) \\
&= \int_{s \in S_X} \psi(s) d\sigma_X^0(s), \tag{SI.27}
\end{aligned}$$

which completes the proof. \square

While Proposition 1 applies to discounted games with $\lambda < 1$, we can get an analogous statement for undiscounted games by multiplying both sides of Eq. (SI.26) by $1 - \lambda$ and taking the limit $\lambda \rightarrow 1^-$:

Corollary 1. If $\psi : S_X \rightarrow \mathbb{R}$ is a bounded, measurable function, then, when the limit exists,

$$\lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T \int_{(x,y) \in S_X \times S_Y} \left[\psi(x) - \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) = 0 \tag{SI.28}$$

for any memory-one strategy, σ_X , where σ_X^0 is the initial action of player X .

Theorem 2 (Autocratic strategies in arbitrary action spaces). Suppose that $\sigma_X[x, y]$ is a memory-one strategy for player X and let σ_X^0 be player X 's initial action. If, for some bounded function, ψ , the equation

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \tag{SI.29}$$

holds for each $x \in S_X$ and $y \in S_Y$, then σ_X^0 and $\sigma_X[x, y]$ together enforce the linear payoff relationship

$$\alpha \pi_X + \beta \pi_Y + \gamma = 0 \tag{SI.30}$$

for *any* strategy of player Y . In other words, the pair $(\sigma_X^0, \sigma_X[x, y])$ is an autocratic strategy.

Proof. If Eq. (SI.29) holds, then by Eq. (SI.26) in Proposition 1,

$$\begin{aligned}
& \alpha \pi_X + \beta \pi_Y + \gamma + (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \\
&= (1 - \lambda) \sum_{t=0}^{\infty} \lambda^t \int_{(x,y) \in S_X \times S_Y} \left[\psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) \right] d\nu_t(x, y) \\
&= (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s), \tag{SI.31}
\end{aligned}$$

and it follows at once that $\alpha\pi_X + \beta\pi_Y + \gamma = 0$. \square

SI.2.1. Two-point autocratic strategies.

Corollary 2. Let $\alpha, \beta, \gamma \in \mathbb{R}$ and suppose that there exist $s_1, s_2 \in S_X$ (i.e. two discrete actions) and $\phi > 0$ such that

$$-\frac{1 - (1 - \lambda)p_0}{\phi} \leq \alpha u_X(s_1, y) + \beta u_Y(s_1, y) + \gamma \leq -\frac{(1 - \lambda)(1 - p_0)}{\phi}; \quad (\text{SI.32a})$$

$$\frac{(1 - \lambda)p_0}{\phi} \leq \alpha u_X(s_2, y) + \beta u_Y(s_2, y) + \gamma \leq \frac{\lambda + (1 - \lambda)p_0}{\phi} \quad (\text{SI.32b})$$

for each $y \in S_Y$, where p_0 is the probability that X initially uses s_1 and $1 - p_0$ is the probability that X initially uses s_2 . Let δ_s be the Dirac measure on S centered at s , and, for $x \in \{s_1, s_2\}$ and $y \in S_Y$, consider the memory-one strategy,

$$\sigma_X[x, y] := p(x, y) \delta_{s_1} + (1 - p(x, y)) \delta_{s_2}, \quad (\text{SI.33})$$

where

$$p(s_1, y) := \frac{1}{\lambda} \left(\phi (\alpha u_X(s_1, y) + \beta u_Y(s_1, y) + \gamma) - (1 - \lambda)p_0 + 1 \right); \quad (\text{SI.34a})$$

$$p(s_2, y) := \frac{1}{\lambda} \left(\phi (\alpha u_X(s_2, y) + \beta u_Y(s_2, y) + \gamma) - (1 - \lambda)p_0 \right). \quad (\text{SI.34b})$$

Then, irrespective of player Y 's strategy, this choice of p_0 and $\sigma_X[x, y]$ enforces $\alpha\pi_X + \beta\pi_Y + \gamma = 0$.

Proof. By Theorem 1, we need only show that there exists $\psi : \{s_1, s_2\} \rightarrow \mathbb{R}$ such that

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \quad (\text{SI.35})$$

for each $x \in S_X$ and $y \in S_Y$ in order to establish the equation $\alpha\pi_X + \beta\pi_Y + \gamma = 0$. Indeed, since we are restricting X 's actions to two points, we may assume that $S_X = \{s_1, s_2\}$. Fix $\psi(s_1) := \psi_1 \in \mathbb{R}$ and let $\psi(s_2) := \frac{1}{\phi} + \psi_1$. Since

$$\begin{aligned} \psi(s_1) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[s_1, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \\ = \psi_1 - \lambda \left(\psi_1 + (1 - p(s_1, y)) \frac{1}{\phi} \right) - (1 - \lambda) \left(\psi_1 + (1 - p_0) \frac{1}{\phi} \right) \\ = \alpha u_X(s_1, y) + \beta u_Y(s_1, y) + \gamma \end{aligned} \quad (\text{SI.36})$$

and

$$\begin{aligned} \psi(s_2) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[s_2, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \\ = \frac{1}{\phi} + \psi_1 - \lambda \left(\psi_1 + (1 - p(s_2, y)) \frac{1}{\phi} \right) - (1 - \lambda) \left(\psi_1 + (1 - p_0) \frac{1}{\phi} \right) \\ = \alpha u_X(s_2, y) + \beta u_Y(s_2, y) + \gamma, \end{aligned} \quad (\text{SI.37})$$

and since $0 \leq p(x, y) \leq 1$ for each $x \in \{s_1, s_2\}$ and $y \in S_Y$ by Eq. (SI.32), the proof is complete. \square

Remark 2. In the undiscounted case ($\lambda = 1$), Eq. (SI.32) is satisfied for some $\phi > 0$ if and only if there exist $s_1, s_2 \in S_X$ such that

$$\alpha u_X(s_1, y) + \beta u_Y(s_1, y) + \gamma \leq 0 \leq \alpha u_X(s_2, y) + \beta u_Y(s_2, y) + \gamma \quad (\text{SI.38})$$

for every $y \in S_Y$. Moreover, if Eq. (SI.32) holds for some λ, p_0, ϕ , and $s_1, s_2 \in S_X$, then it must be true that Eq. (SI.38) also holds for every $y \in S_Y$. If Eq. (SI.38) does not hold for a particular choice of $s_1, s_2 \in S_X$, then s_1 and s_2 cannot form a two-point autocratic strategy for any discounting factor, λ . Therefore, Eq. (SI.38), which is easy to check, offers a straightforward way to show that two actions cannot form a two-point autocratic strategy for a particular game.

Remark 3. For ψ , s_1 , and s_2 , fixed, one can ask which strategies of the form

$$\sigma_X^0 = p_0 \delta_{s_1} + (1 - p_0) \delta_{s_2}; \quad (\text{SI.39a})$$

$$\sigma_X[x, y] = p(x, y) \delta_{s_1} + (1 - p(x, y)) \delta_{s_2}, \quad (\text{SI.39b})$$

for some p_0 and $p(x, y)$, satisfy the equation

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s). \quad (\text{SI.40})$$

Indeed, we see from the proof of Corollary 2 that, for a strategy of this form, we must have

$$p(x, y) = \frac{\psi(s_2) - \frac{1}{\lambda} \left(\psi(x) - \left(\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma + (1 - \lambda) \left(\psi(s_1) p_0 + \psi(s_2) (1 - p_0) \right) \right) \right)}{\psi(s_2) - \psi(s_1)} \quad (\text{SI.41})$$

for each $x \in \{s_1, s_2\}$ and $y \in S_Y$. Therefore, this simple case does not capture the generally-implicit nature of autocratic strategies because one can explicitly write down two-point strategies via Eq. (SI.41), which is typically not possible for strategies concentrated on more than just two actions.

SI.3. EXAMPLES

Here we present some simple examples of Theorem 1 and its implications. In §SI.3.1, we demonstrate how Theorem 1 reduces to the main result of Press and Dyson (2012) when the action space has only two options. Moreover, we use an action space consisting of three choices to illustrate the implicit nature of autocratic strategies defined via Press-Dyson functions for more than two actions. In §SI.3.2, we show that there is no way for a player to unilaterally set her own score using Theorem 1. In particular, despite the implicit nature of autocratic strategies, one can use Theorem 1 to deduce the non-existence of certain classes of strategies.

SI.3.1. Games with finitely many actions and no discounting. Suppose that $S_X = S_Y = \{A_1, \dots, A_n\}$ is finite. If $\psi : S_X \rightarrow \mathbb{R}$ and σ_X is a memory-one strategy, then, for $\lambda = 1$,

$$\begin{aligned} \psi(x) - \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) &= \psi(x) - \sum_{r=1}^n \psi(A_r) \sigma_X[x, y](A_r) \\ &= \psi(x) - \sum_{r=1}^{n-1} \psi(A_r) \sigma_X[x, y](A_r) - \psi(A_n) \left(1 - \sum_{r=1}^{n-1} \sigma_X[x, y](A_r) \right) \\ &= \psi(x) - \psi(A_n) - \sum_{r=1}^{n-1} \left(\psi(A_r) - \psi(A_n) \right) \sigma_X[x, y](A_r) \\ &= \begin{cases} \sum_{r=1}^{n-1} \left(\psi(A_n) - \psi(A_r) \right) \left(\sigma_X[x, y](A_r) - \delta_{r,r'} \right) & x = A_{r'} \neq A_n, \\ \sum_{r=1}^{n-1} \left(\psi(A_n) - \psi(A_r) \right) \sigma_X[x, y](A_r) & x = A_n. \end{cases} \end{aligned} \quad (\text{SI.42})$$

Therefore, if $n = 2$ and $S_X = S_Y = \{A_1, A_2\}$, we have

$$\psi(x) - \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) = \left(\psi(A_2) - \psi(A_1) \right) \tilde{\mathbf{p}}. \quad (\text{SI.43})$$

Thus, a Press-Dyson function is a scalar multiple of $\tilde{\mathbf{p}}$ (defined by Eq. (2) in the main text).

On the other hand, if $n = 3$ and $S_X = S_Y = \{A_1, A_2, A_3\}$, then

$$\psi(x) - \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) = \begin{cases} c_1 \sigma_X[x, y](A_1) + c_2 \sigma_X[x, y](A_2) - c_1 & x = A_1, \\ c_1 \sigma_X[x, y](A_1) + c_2 \sigma_X[x, y](A_2) - c_2 & x = A_2, \\ c_1 \sigma_X[x, y](A_1) + c_2 \sigma_X[x, y](A_2) & x = A_3, \end{cases} \quad (\text{SI.44})$$

where $c_1 = \psi(A_3) - \psi(A_1)$ and $c_2 = \psi(A_3) - \psi(A_2)$. For each $x \in S_X$ and $y \in S_Y$, the measure $\sigma_X[x, y]$ is uniquely determined by $\sigma_X[x, y](A_1)$ and $\sigma_X[x, y](A_2)$. Therefore, unlike for $n = 2$, one cannot necessarily eliminate both of c_1 and c_2 by normalizing the Press-Dyson function. For this reason, as well as due to the fact that a Press-Dyson function reduces to a multiple of the Press-Dyson vector for $n = 2$, it is perhaps more natural to refer to any multiple of $\tilde{\mathbf{p}}$ (as opposed to just $\tilde{\mathbf{p}}$ itself) as a Press-Dyson vector.

SI.3.2. Continuous Donation Game. Here we establish some further results claimed in the main text for extortionate, generous, and equalizer strategies for the continuous Donation Game.

SI.3.2.1. Relationship to the classical Donation Game. For a two-point strategy, X 's action space may be restricted to $S_X = \{0, K\}$. Therefore, the scaling function $\psi : S_X \rightarrow \mathbb{R}$ of Theorem 1 is defined by two numbers, $\psi(0)$ and $\psi(K)$. Letting $\phi := 1/(\psi(0) - \psi(K))$, we see by Corollary 2 in §SI.2 that the function

$$p(x, y) := \begin{cases} \frac{1}{\lambda} \left(1 - \phi \left(\chi b(K) + c(K) - b(y) - \chi c(y) - (\chi - 1) \kappa \right) - (1 - \lambda) p_0 \right) & x = K, \\ \frac{1}{\lambda} \left(\phi \left(b(y) + \chi c(y) + (\chi - 1) \kappa \right) - (1 - \lambda) p_0 \right) & x = 0 \end{cases} \quad (\text{SI.45})$$

gives well-defined reaction probabilities provided ψ (hence ϕ) and p_0 are chosen so that $0 \leq p(x, y) \leq 1$ for each $x \in \{0, K\}$ and $y \in [0, K]$. For any such ψ , the memory-one strategy,

$$\sigma_X[x, y] := \left(1 - p(x, y) \right) \delta_0 + p(x, y) \delta_K, \quad (\text{SI.46})$$

together with $\sigma_X^0 = (1 - p_0) \delta_0 + p_0 \delta_K$, defines an autocratic strategy that allows X to enforce the linear relationship

$$\pi_X - \kappa = \chi (\pi_Y - \kappa). \quad (\text{SI.47})$$

Note Eq. (SI.46) simply states formally that player X fully cooperates with probability $p(x, y)$ and defects with probability $1 - p(x, y)$ following an outcome in which X plays x and Y plays y .

In the absence of discounting, i.e. in the limit $\lambda \rightarrow 1$, we have

$$p(K, K) = 1 - \phi(\chi - 1) \left(b(K) - c(K) - \kappa \right); \quad (\text{SI.48a})$$

$$p(K, 0) = 1 - \phi \left(\chi b(K) + c(K) - (\chi - 1) \kappa \right); \quad (\text{SI.48b})$$

$$p(0, K) = \phi \left(b(K) + \chi c(K) + (\chi - 1) \kappa \right); \quad (\text{SI.48c})$$

$$p(0, 0) = \phi(\chi - 1) \kappa. \quad (\text{SI.48d})$$

Thus, setting $b := b(K)$, and $c := c(K)$, the general form, Eq. (SI.48), recovers the discrete-action-space case, Eq. (6). In particular, the autocratic memory-one strategy in Eq. (SI.46) is a direct generalization of zero-determinant strategies to the continuous Donation Game. However, this strategy contains much more information than the corresponding strategy for the classical Donation Game because it encodes X 's play in response to Y 's for every $y \in [0, K]$. Despite the fact that player Y has an uncountably infinite number of actions to choose from, player X can still ensure that Eq. (SI.47) holds by playing only two actions.

SI.3.2.2. Extortionate and generous strategies. In the main text, we saw that X can unilaterally enforce $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ for $\chi \geq 1$ provided $0 \leq \kappa \leq b(K) - c(K)$ and λ is sufficiently close to 1. If $\chi = 1$, then κ is irrelevant and the linear relationship is simply $\pi_X = \pi_Y$. Here, we show that, if $\chi > 1$, then




FIGURE 4. Extortion in the continuous Donation Game without discounting. (A) shows linear costs, $c(s) = 2s$, and saturating benefits, $b(c(s)) = 5(1 - e^{-c(s)})$. In (B), the reaction function of player X , $p(x, y) = p(y)$ (see Eq. (8)), indicates the probability that player X plays K (as opposed to 0) after Y uses y and represents a reactive strategy because it depends solely on Y 's previous move. In (C), the function $\psi(s) = -\chi b(s) - c(s)$ is shown together with two dashed lines indicating $\psi(x)$ and $\int_{s \in S_X} \psi(s) d\sigma_X[x, y](s)$, respectively, for $x = 0.3$ where $\sigma_X[x, y]$ denotes the memory-one strategy for player X based on the reactive strategy in (A). The vertical distance between the dashed lines must equal $\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = u_X(x, y) - \chi u_Y(x, y)$ for each $x, y \in [0, 2]$ to satisfy Theorem 1. Parameters: $\lambda = 1, \kappa = 0, \alpha = 1, \beta = -2$, and $\gamma = 0$, yielding an extortion factor of $\chi := -\beta = 2$.

$0 \leq \kappa \leq b(K) - c(K)$ is necessary for such a payoff relationship to be enforced via Theorem 1. Indeed, if

$$\begin{aligned} (u_X(x, y) - \kappa) - \chi(u_Y(x, y) - \kappa) &= -\chi b(x) - c(x) + b(y) + \chi c(y) + (\chi - 1)\kappa \\ &= \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \end{aligned} \quad (\text{SI.49})$$

for some bounded $\psi : S_X \rightarrow \mathbb{R}$ and each $x \in S_X$ and $y \in S_Y$, then we obtain

$$(\chi - 1) \geq \sup \psi - \lambda \sup \psi - (1 - \lambda) \sup \psi = 0 \quad (\text{SI.50})$$

by taking $y = 0$ and letting $\psi(x)$ approach $\sup \psi$, and we see that

$$-\chi b(K) - c(K) + b(K) + \chi c(K) + (\chi - 1)\kappa \leq \inf \psi - \lambda \inf \psi - (1 - \lambda) \inf \psi = 0 \quad (\text{SI.51})$$

by taking $y = K$ and letting $\psi(x)$ approach $\inf \psi$. Since $\chi > 1$, we have $0 \leq \kappa \leq b(K) - c(K)$.

SI.3.2.3. Opponent-equalizing strategies. Although a player cannot set her own score in the continuous Donation Game, she can set the score of her opponent. We saw in the main text that X can set Y 's score to anything between 0 and $b(K) - c(K)$ provided λ is sufficiently large, and here we show that this interval is the only range of payoffs for player Y that X can unilaterally set via Theorem 1. Indeed, if γ satisfies

$$\begin{aligned} u_Y(x, y) - \gamma &= b(x) - c(y) - \gamma \\ &= \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \end{aligned} \quad (\text{SI.52})$$

for some bounded $\psi : S_X \rightarrow \mathbb{R}$ and each $x \in S_X$ and $y \in S_Y$, then we obtain

$$-\gamma \leq \inf \psi - \lambda \inf \psi - (1 - \lambda) \inf \psi = 0 \quad (\text{SI.53})$$

by taking $y = 0$ and letting $\psi(x)$ approach $\inf \psi$, and we find that

$$b(K) - c(K) - \gamma \geq \sup \psi - \lambda \sup \psi - (1 - \lambda) \sup \psi = 0 \quad (\text{SI.54})$$

by taking $y = K$ and letting $\psi(x)$ approach $\sup \psi$. It follows immediately that $0 \leq \gamma \leq b(K) - c(K)$.

SI.3.2.4. *Self-equalizing strategies.* We saw in the main text that extortionate strategies exist in the continuous Donation Game, as demonstrated by the two-point strategy defined by Eq. (8). However, it certainly need not be the case that for any $\alpha, \beta, \gamma \in \mathbb{R}$, there exists $\psi : S_X \rightarrow \mathbb{R}$ such that

$$\alpha u_X(x, y) + \beta u_Y(x, y) + \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \quad (\text{SI.55})$$

for some $(\sigma_X^0, \sigma_X[x, y])$ and each $x \in S_X$ and $y \in S_Y$. For example suppose that $\alpha = 1$, $\beta = 0$, and γ is some fixed real number. For $u_X(x, y) = b(y) - c(x)$, the equation

$$b(y) - c(x) - \gamma = u_X(x, y) - \gamma = \psi(x) - \lambda \int_{s \in S_X} \psi(s) d\sigma_X[x, y](s) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \quad (\text{SI.56})$$

holds for some $\psi : S_X \rightarrow \mathbb{R}$ and $(\sigma_X^0, \sigma_X[x, y])$, only if

$$\lambda \inf \psi \leq \psi(x) + c(x) + \gamma - b(y) - (1 - \lambda) \int_{s \in S_X} \psi(s) d\sigma_X^0(s) \leq \lambda \sup \psi \quad (\text{SI.57})$$

for each $x, y \in [0, K]$. However, these inequalities imply that

$$(1 - \lambda) \inf \psi \geq (1 - \lambda) \sup \psi + (b(K) - c(K)), \quad (\text{SI.58})$$

which is impossible since $b(K) > c(K)$ and $\inf \psi \leq \sup \psi$. Thus, there is no feasible Press-Dyson function that allows player X to unilaterally enforce the equation $\pi_X = \gamma$. In other words, a player cannot unilaterally set her own payoff.

SI.3.2.5. *Initial actions.* Here we state the conditions on p_0 (for two-point strategies) and x_0 (for deterministic strategies) that allow X to enforce autocratic strategies in the continuous Donation Game.

If p_0 denotes the probability that X initially plays $x = K$ (as opposed to $x = 0$), then the function

$$p(x, y) = \frac{1}{\lambda} \left(\frac{b(y) + \chi c(y) + (\chi - 1)\kappa - (1 - \lambda)(\chi b(K) + c(K))p_0}{\chi b(K) + c(K)} \right) \quad (\text{SI.59})$$

gives a well-defined reaction probability provided λ satisfies Eq. (7) in the main text and p_0 satisfies

$$\begin{aligned} & \max \left\{ \frac{(\chi - 1)\kappa - \lambda(\chi b(K) + c(K)) + (b(K) + \chi c(K))}{(1 - \lambda)(\chi b(K) + c(K))}, 0 \right\} \\ & \leq p_0 \leq \min \left\{ \frac{(\chi - 1)\kappa}{(1 - \lambda)(\chi b(K) + c(K))}, 1 \right\}. \end{aligned} \quad (\text{SI.60})$$

Moreover, for such a p_0 and $p(x, y)$, the two-point strategy defined by

$$\sigma_X^0 = (1 - p_0)\delta_0 + p_0\delta_K; \quad (\text{SI.61a})$$

$$\sigma_X[x, y] = (1 - p(x, y))\delta_0 + p(x, y)\delta_K \quad (\text{SI.61b})$$

allows X to enforce $\pi_X - \kappa = \chi(\pi_Y - \kappa)$. This reaction probability was obtained using the scaling function $\psi(s) := -\chi b(s) - c(s)$ on $S_X = \{0, K\}$.

If X wishes to enforce $\pi_Y = \gamma$ for $0 \leq \gamma \leq b(K) - c(K)$, she may use $\psi(s) := b(s)$ on $S_X = \{0, K\}$ and

$$p(x, y) = \frac{1}{\lambda} \left(\frac{c(y) + \gamma}{b(K)} - (1 - \lambda)p_0 \right), \quad (\text{SI.62})$$

provided $\lambda \geq c(K)/b(K)$ and

$$\max \left\{ \frac{\gamma - \lambda b(K) + c(K)}{(1 - \lambda)b(K)}, 0 \right\} \leq p_0 \leq \min \left\{ \frac{\gamma}{(1 - \lambda)b(K)}, 1 \right\}. \quad (\text{SI.63})$$

Using $\psi(s) := -\chi b(s) - c(s)$ on $S_X = [0, K]$, X can also enforce $\pi_X - \kappa = \chi(\pi_Y - \kappa)$ using the deterministic strategy with reaction function

$$r^X(x, y) = \psi^{-1} \left(\frac{-b(y) - \chi c(y) - (\chi - 1)\kappa - (1 - \lambda)\psi(x_0)}{\lambda} \right), \quad (\text{SI.64})$$

where $\psi^{-1}(\dots)$ denotes the inverse of the function ψ , provided Eq. (7) holds and X 's initial action, x_0 , satisfies

$$\frac{(\chi - 1)\kappa - \lambda(\chi b(K) + c(K)) + (b(K) + \chi c(K))}{1 - \lambda} \leq \chi b(x_0) + c(x_0) \leq \frac{(\chi - 1)\kappa}{1 - \lambda}. \quad (\text{SI.65})$$

For X to set $\pi_Y = \gamma$ for $0 \leq \gamma \leq b(K) - c(K)$ using a deterministic strategy, then she can do so by using $\psi(s) := b(s)$ and the reaction function

$$r^X(x, y) = \psi^{-1} \left(\frac{c(y) + \gamma - (1 - \lambda)\psi(x_0)}{\lambda} \right), \quad (\text{SI.66})$$

provided $\lambda \geq c(K)/b(K)$ and X 's initial action, x_0 , satisfies

$$\frac{\gamma - \lambda b(K) + c(K)}{1 - \lambda} \leq b(x_0) \leq \frac{\gamma}{1 - \lambda}. \quad (\text{SI.67})$$